

CLAIMS

1. A switching system, comprising:
 - 5 a plurality of virtual queue managers that store data; and
 - a plurality of switch circuits, each of said switch circuits being operatively connected to each of said virtual queue managers, and at least one of said switch circuits having an internal scheduler,
 - 10 wherein said internal scheduler selects at least one of said virtual queue managers to send data to said plurality of switch circuits.
 - .
2. A switching system as recited in claim 1, wherein in a given period of time, a cell of data from one of said virtual queue managers is concurrently switched through said switch circuits.
 - 15
3. A switching system as recited in claim 1, wherein the selected at least one of said virtual queue managers sends a cell of data by sending a block of data to each of said plurality of switches, wherein each of the blocks includes a header and payload data.
 - 20
4. A switching system as recited in claim 1,
 - wherein said virtual queue manager stores cells, and
 - wherein said virtual queue manager transfers one of the cells to said switch circuits by concurrently directing each of a plurality of blocks
 - 25 representing the cell to a different ones of said switch circuits.
5. A switching system as recited in claim 4, wherein each of the blocks include a header, a payload header and payload data.

6. A switching system as recited in claim 5, wherein each of said switching circuits, upon receiving one of the blocks, operates to switch the received one of the blocks in accordance with the payload header.
- 5 7. A switching system as recited in claim 1, wherein each of said switching circuits comprise a concurrent switch.
8. A switching system as recited in claim 1, wherein the data is stored as cells in said virtual queue managers, and
- 10 wherein each of the cells include a header, a payload header and payload data.
9. A switching system as recited in claim 1, wherein said internal scheduler controls said virtual queue managers such that data is transferred
- 15 to said plurality of switch circuits from said virtual queue managers under the control of said internal scheduler.
10. A switching system as recited in claim 9, wherein said internal scheduler does not directly couple to any other of said switch chips other than
- 20 the at least one of said switch circuits having an internal scheduler.
11. A switching system as recited in claim 1, wherein said switch circuits do not provide any direct communication links between said internal scheduler and any other of said switch circuits other than the at least one of
- 25 said switch circuits having an internal scheduler.
12. A multi-port, pipelined concurrent switching apparatus, comprising a switch fabric that can aggregate links and multiple switching chips without inter-chip scheduling communications between any of said multiple switching
- 30 chips, thereby supporting higher switching bandwidth per port.

13. A multi-port, pipelined concurrent switching apparatus as recited in claim 12, wherein said switch fabric operates in accordance with a pipeline when performing switching operations.

5

14. A multi-port, pipelined concurrent switching apparatus as recited in claim 13, wherein the pipeline includes a request stage, a grant stage, a data-in transfer, and a data-out transfer, with the request stage being the initial stage of the pipeline and the data-out transfer being the final stage of the
10 pipeline.

15. A switching apparatus, comprising:

a plurality of queues for storing blocks of data, each block including a request and a payload, and each of the requests stored in said queues being associated with one of the payloads stored in said queues; and

a plurality of switches, at least one of said switches including a scheduler that arbitrates requests sent by said queues to said scheduler,

wherein, when issued to the scheduler, each of said requests operates to request switching the associated payload through said switches in a
20 particular manner, and

wherein each of the requests are issued in advance of sending the associated payloads, and only after a particular request is granted does the associated payload get transmitted from the associated one or more of said queues to said switches where the associated payload is passed through said
25 switches.

16. A method for operating a switching apparatus, said method comprising:

(a) sending a first request to transmit a first payload through the
30 switching apparatus;

- 6
5
4
3
2
1
- (b) determining whether the first request has been granted;
 - (c) sending a subsequent request to transmit a subsequent payload through the switching apparatus; and
 - (d) sending, concurrently with said sending (c), the first payload when said determining (d) has determined that the first grant has been granted.

17. A method as recited in claim 16, wherein the switching apparatus includes a plurality of queues and a plurality of switches.

- 10 18. A method as recited in claim 16, wherein at least one of said switches includes a scheduler,
wherein said sending (a) sends the first request to the at least one of said switches that includes a scheduler,
wherein said sending (c) sends the subsequent request to the at least
15 one of said switches that includes a scheduler, and
wherein said sending (d) sends the first payload a plurality of said switches so as to transmit the first payload through said switching apparatus.

- 20 19. A method for operating a switching apparatus having multiple virtual queues and multiple switch chips, said method comprising:
(a) receiving an incoming packet to be passed through the switching apparatus;
(b) dividing the incoming packet into a plurality of fixed length blocks, the blocks including at least a header, a payload and a payload header, the
25 header including control information, and the payload header including a destination indicator for the payload;
(c) temporarily storing the blocks in the virtual queues;

(d) determining when the payloads associated with the blocks associated with the incoming packet that are stored in the virtual queues are to be passed through the switching chips;

- (e) concurrently transferring each of the payloads and their payload headers for the blocks that are associated with the incoming packet from the virtual queues to different ones of the switch chips when said determining (d) determines that the payloads associated with the blocks are to be passed through the switch chips; and
- (f) concurrently switching the blocks associated with the incoming packet through the switch chips in accordance with the destination identifier provided within the payload header of each of the blocks.

20. A method as recited in claim 19, wherein said receiving (a) of the incoming packet is received at an incoming, and wherein the destination identifier indicates a destination port.

21. A method as recited in claim 19, wherein one or more of the fixed length blocks represent a cell, and one or more of the cells represent the incoming packet.

20

22. A method as recited in claim 19, wherein the destination identifier includes a reference indicator for one of the virtual queues.

25

23. A method as recited in claim 19, wherein said determining (d) comprises:

(d1) requesting to transfer the blocks associated with the incoming packet; and

(d2) receiving a grant to transfer the cells associated with the incoming packet.

30

24. A method as recited in claim 23, wherein the grant informs the one or of the virtual queues of the blocks to be concurrently transferred during said transferring (e).

5 25. A method as recited in claim 19, wherein at least one of said switch chips comprises a scheduler that determines when blocks stored within the virtual queues are passed through the switching chips.

10 26. A method as recited in claim 25, wherein said determining (d) comprises:

(d1) sending requests to transfer blocks from one or more of the virtual queues to the scheduler, and

(d2) determining, at the scheduler, one or more of the requests to be simultaneously granted; and

15 (d3) informing the one or more of the virtual queues of the one or more of the requests that have been granted.

27. A method as recited in claim 26, wherein said sending (d1) sends subsequent requests to transfer blocks from one or more of the virtual queues 20 concurrently with said transferring (e) of blocks associated with an earlier request.

28. A method as recited in claim 27, wherein the requests are provided within at least one of the blocks associated with the incoming packet.

25

29. A method as recited in claim 26, wherein the requests are provided within at least one of the blocks associated with the incoming packet.

30. A method as recited in claim 25, wherein said dividing (b) operates to divide the incoming packet into blocks of multiple types.

31. A method as recited in claim 30, wherein the header of each of the 5 blocks further includes a block type so that the multiple types of blocks can be distinguished.

32. A method as recited in claim 30, wherein the multiple types includes at least a first type and a second type, wherein the size of the header of the 10 blocks of the second type is smaller than the size of the header of the blocks of the first type, and wherein the amount of the data within the blocks of the second type is greater than the amount of the data within the blocks of the first type.

15 33. A method as recited in claim 32, wherein the blocks of the first type are sent to the switching chip including the scheduler, and the blocks of the second type are sent to the switch chips not including a scheduler.

20 34. A method as recited in claim 33, wherein the blocks of the first type can include a request.

35. A method for switching a block of data through a switch system having virtual queues and switching devices, at least one of the switching devices including a scheduler, said method comprising:

25 (a) receiving, at a switching device including a scheduler, a block from a virtual queue, the block including a header, a payload and a payload header;

(b) directing the header of the block to the scheduler, and directing the payload and the payload header to the switching device; and

(c) switching the payload through the switching device in accordance with the payload header.

36. A method as recited in claim 35,

5 wherein the payload header includes a destination header, and

wherein the switching device is a concurrent switch, and wherein said switching (c) operates to configure the concurrent switch in accordance with the destination identifier provided in the payload header.

10 37. A method as recited in claim 36, wherein the payload and the payload header of the block pertain to an earlier request that has been granted by the scheduler, and wherein the header of the block pertains to a later request.

15 38. A method as recited in claim 35, wherein the payload and the payload header of the block pertain to an earlier request that has been granted by the scheduler, and wherein the header of the block pertains to a later request.

39. A method as recited in claim 35, wherein the switching device is a concurrent switch.

20

40. A switching system, comprising:

a plurality of virtual queue managers that store data; and

25 a plurality of switch circuits, each of said switch circuits being operatively connected to each of said virtual queue managers, and each of said switch circuits including at least an internal scheduler,

wherein said internal schedulers operate to switch in a synchronized manner to thereby concurrently send data, provided by one or more of said virtual queue managers, through said plurality of switch circuits.

41. A switching system as recited in claim 40, wherein in a given period of time, a cell of data from one of said virtual queue managers is concurrently switched through said switch circuits.

5 42. A switching system as recited in claim 40, wherein said one or more virtual queue managers send a cell of data by sending a block of data to each of said plurality of switches, wherein each of the blocks includes a header and payload data.

10 43. A switching system as recited in claim 40,
wherein said virtual queue managers store cells, and
wherein said virtual queue managers transfer one of the cells to said switch circuits by concurrently directing each of a plurality of blocks representing the cell to a different ones of said switch circuits.

15 44. A switching system as recited in claim 43, wherein each of the blocks include a header and payload data.

20 45. A switching system as recited in claim 44, wherein each of said switching circuits, upon receiving one of the blocks, operates to switch the received one of the blocks in accordance with switch information provided by said internal scheduler provided therein.

25 46. A switching system as recited in claim 40, wherein each of said switching circuits comprise a concurrent switch.

47. A switching system as recited in claim 40, wherein said internal scheduler does not directly couple to any other of said switch chips other than said switch circuits in which it is internal.

48. A switching system as recited in claim 40, wherein said switch circuits do not provide any direct communication links between and any other of said switch circuits.

5

49. A method for operating a switching apparatus having multiple virtual queues and multiple switch chips, said method comprising:

- (a) receiving an incoming packet to be passed through the switching apparatus;
- 10 (b) dividing the incoming packet into a plurality of fixed length blocks, the blocks including at least a header and a payload, the header including control information;
- (c) temporarily storing the blocks in the virtual queues;
- (d) concurrently transferring the header and the payloads for the 15 blocks that are associated with the incoming packet from the virtual queues to different ones of the switch chips;
- (e) scheduling, independently and simultaneously at each of the switch chips, when the payloads associated with the blocks associated with the incoming packet are to be passed through the switching chips; and
- 20 (f) concurrently switching the blocks in accordance with said scheduling (e).

50. A method for switching a block of data through a switch system having a plurality of virtual queues and a plurality of switching devices, each of the 25 switching devices including an internal scheduler and a switch, said method comprising:

- (a) receiving, at each of the switching devices, a header from one or more of the virtual queues and at least a portion of a payload from one or more of the virtual queues;

(b) directing, at each of the switching device, the header to the internal scheduler and directing at least a portion of the payload to the switch;

(c) producing, at each of the schedulers, switching information based on the header from one or more of the virtual queues; and

5 (d) switching the payload through the switch in accordance with the switching information.

51. A method as recited in claim 50, wherein the switching device is a concurrent switch, and

10 wherein said switching (d) operates to configure the concurrent switch in accordance with the switching information.

52. A method as recited in claim 50, wherein said directing (b) causes each of the switching devices to each the headers from each of the virtual
15 queues.

53. A method as recited in claim 50, wherein said producing (c) operates to separately and independently produce the switch information at each of the schedulers.

20

54. A method for switching a block of data through a switch system having a plurality of packet forwarding devices and a plurality of switch devices, the switch devices include at least virtual queues and switch circuits, said method comprising:

25 (a) receiving, at each of the switching devices, a header and at least a portion of a payload from one or more of the packet forwarding devices;

(b) directing, at each of the switching devices, the payload to the virtual queues based on the header;

(c) storing the payload in the virtual queues; and

(d) switching the payload from a source location in the virtual queues to a destination location in the virtual queues based on determined control information.

5 55. A method as recited in claim 54, wherein the packet forwarding devices duplicate the headers such that they are supplied to each of the switching devices.

10 56. A method as recited in claim 54, wherein the packet forwarding devices are traffic managers or network processors.

57. A method as recited in claim 54, wherein the switch devices use on-chip or off-chip memory to implement the virtual queues.

15 58. A method as recited in claim 54, wherein the switching devices operate to independently and synchronously.

59. A method as recited in claim 54, wherein the control information is determined from the headers.

20

60. A method as recited in claim 59, wherein the control information is a destination identifier.